To appear in *Philosophical Topics* special issue on Perceptual Appearances, (eds) C. Hill and B. McLaughlin.

Naïve Realism and the Science of Illusion

IAN PHILLIPS

ABSTRACT: Critics have long complained that naïve realism cannot adequately account for perceptual illusion. This complaint has a tendency to ally itself with the aspersion that naïve realism is hopelessly out of touch with vision science. Here I offer a partial reply to both complaint and aspersion. I do so by showing how careful reflection on a simple, empirically grounded model of illusion reveals heterodox ways of thinking about familiar illusions which are quite congenial to the naïve realist.

1. Naïve Realism and Illusions

As Martin characterizes it, naïve realism is the view that "the actual objects of perception, the external things such as trees, tables and rainbows, which one can perceive, and the properties which they can manifest to one when perceived, partly constitute one's conscious experience, and hence determine the phenomenal character of one's experience" (2009 [1997], p. 93). Martin emphasizes that this talk of constitution and determination is intended quite literally, highlighting the consequence "that one could not be having the very experience one has, were the objects perceived not to exist, or were they to lack the features they are perceived to have" (ibid.).[1]

So characterized, it may seem obvious that naïve realism cannot be extended to cases of illusion. As a result, illusions form the basis of a long-standing critique of the view.[2]

[1] Recent defences of naïve realism in various guises include: Martin 2002, 2004 and 2006, Travis 2004, Campbell 2002 and 2009, and Brewer 2011 and 2013.

[2] For recent variants on this old objection see: Foster 2000, Burge 2005, Byrne 2009, Smith 2010, Block 2010, McLaughlin 2010, and Millar 2015. Siegel 2011 contains closely related criticism.

Illusions, so the thought goes, involve a subject misperceiving an object as being some way it is not: as bent when straight, as yellow when green, as concave when convex. The character of our experience in such cases cannot then be explained by appeal to properties which the object actually makes manifest. For instance, Smith asks us to imagine an illusion in which a green square appears yellow, perhaps due to peculiar lighting (2010, p. 392). "How could a green square's being a constituent of my experience account for the apparent yellowness of the square?" he challenges. "By itself, it clearly cannot." (p. 389)

Martin appears to accept this verdict, at least regarding certain cases of illusion, acknowledging that the "naïve realist account of perceptual experience … cannot be directly applied to any case of delusive experience, such as illusions where one does perceive an external object, but misperceives it as other than it really is" (2009 [1997], p. 95). In response, he proposes a disjunctivist treatment of perceptual experience which conceives of veridical experience as fundamentally different in kind to both illusions and hallucinations.[3] Against this, critics of naïve realism have argued that, even if a disjunctivist account of hallucination is tenable, it cannot be extended to illusions. Foster (2000, part II), for example, appeals to a sequence of cases in which a circular shape is seen through a series of lenses, the first highly distorting, the last not distorting at all. He argues that there are "no sharp discontinuities" across the series despite the fact that only the final case is veridical. "[I]t is surely clear", Foster concludes, "that all the perceptions in the series have to be thought of as of the same general kind" (p. 69). Smith (2010) offers a slightly different argument, complaining that the naïve realist cannot accommodate the *partiality* of illusion, for instance the fact that we can correctly see a green square as square even though we misperceive its colour. Neither argument is unproblematic. Each relies on controversial background assumptions whose status the naïve realist will likely wish to question. To pick just two: Foster's argument requires that only fully determinate aspects of the external world can be appealed to in

[3] One way of developing such a position would be to treat the relevant illusions on lines analogous to Martin's approach to object hallucination (see his 2004 and 2006). These illusions would then be characterizable as episodes whose fundamental nature was partly specifiable only in epistemic terms: as episodes which are in some relevant respect indiscriminable upon reflection from a veridical perceptual experience of a certain kind (cf. Martin 2004, p. 81 and 2002, p. 395, fn. 24).

accounting for phenomenal character (2000, p. 71); and Smith's that "visually perceiving an object's shape requires seeing that object's colour" (2010, p. 404).[4]

However, even if no decisive argument is forthcoming, critics of naïve realism are often content to present themselves as offering a *better* explanation of illusion than the naïve realist. These critics typically hold that experience in general is representational. Given this, illusions are straightforwardly construed as occasions on which experience *misrepresents* its objects. After all, it is in the nature of representation, in contrast to presentation, that features can be represented despite being non-actual. In this way, representationalists cite illusions as powerful abductive support for their view. Thus, Tye:

> In cases of illusion, the perceived object appears other than it is. In such cases, intuitively, the perceptual experience is inaccurate. And it is so precisely because the object is not as it appears to be. The simplest explanation of this, in my view, is that, where there is a perceived object, a perceptual experience has a content into which the perceived object enters along with its apparent properties. (2014, p. 293)

In turn this suggests that unless the naïve realist can provide a satisfactory alternative account of illusions—one on an equal or better footing than the representationalist's—we should prefer the "simple" view offered by the representationalist. In this way, McLaughlin, criticizing Campbell's version of naïve realism, complains: "one may be seeing a scene, and yet the scene looks some way that it isn't. To accommodate this, it seems that we have to

[4] Fish (2009, pp. 43-6) and Millar (2015, pp. 612-3) also endorse this principle in making similar arguments against a disjunctivist strategy. Should we? If by colour, Smith means chromatic colour, then shape perception in achromatopsic subjects and in ordinary vision in very low levels of light are plausibly counter-examples. Can Smith's principle be weakened to avoid such counter-examples and yet play the required role in his argument? Suppose that seeing an object's colour is a matter of seeing its hue, saturation and lightness, and that subjects in very low light cannot track variation in hue and saturation but only in lightness. If that is right, such cases may only force us to reformulate Smith's principle as the claim that one cannot visually perceive an object's shape without seeing *some aspect* of its colour, be it hue, saturation or lightness. I am, nonetheless, sceptical that any such reformulation will ultimately prove tenable. But that is matter for a different time. (Note, separately, that the principle arguably needs to be amended to allow for the perception of objects whose shape can be seen through an opaque covering.)

posit that the visual perceptual experience has a representational content. I can find in Campbell no alternative explanation of illusions" (2010, pp. 261-2).

One naïve realist who attempts to meet this objection head-on is Brewer (2008, 2011). The specifics of Brewer's account of illusions have, however, been disputed on empirical grounds. And such criticism fuels the more general fashion of deprecating naïve realism as (to echo the tone of certain critics) hopelessly out of touch with vision science.[5] In what follows I argue that naïve realism is much better placed to avail itself of empirically adequate accounts of illusion than such criticism recognizes. My aim is not to offer "a naïve realist account of illusion", however. I suspect that illusions lack sufficient unity to be treated by any general account. Indeed, some illusions plausibly will require a disjunctivist treatment as Martin suggests above (cf. Brewer 2011, p. 115ff.). Instead, I consider a simple model of certain paradigm illusions in order to bring out ways of thinking about such illusions which are both empirically motivated and congenial to naïve realism. I begin this task in §3. Ahead of that it is helpful briefly to sketch a framework concerning our thought and talk of appearances within which to locate aspects of the discussion to come.

2. Thought and Talk about Looks

[5] See especially Burge 2005, also Searle 2015, p. 165. Burge takes vision science to have established certain very general facts about illusions (viz. his "Proximality Principle") which are supposedly inconsistent with naïve realism. Considered at the level of abstraction at which Burge operates it seems to me that Campbell (2010) offers a sufficient reply on behalf of the naïve realist. Nonetheless one might hanker after a more detailed and local demonstration of how empirical work might mesh with a naïve realist approach to perception. The present paper partly responds to that felt need. Another example of a very general appeal to considerations from vision science in relation to illusions is Antony 2011. Antony hopes to draw from vision science (or more specifically a certain neo-Helmholtzian constructivist approach to vision science found in the works of Gregory, Mack and Palmer) materials which will help make sense of the metaphor of perception as involving an "openness to the world". The difficulty she faces echoes the difficultly which Campbell argues faces Burge. For example, she finds in vision science the thought that "[t]he properties instantiated in perceptual experiences are never direct reflexes of the familiar properties we attribute to objects in the external world" (33) and uses this thought (if I understand her correctly) to place both illusions and veridical perceptions on a par as equally encounters with appearance-properties. What is hard to see, however, is why the claim of *causal* indirectness derived from vision science should bear on the claim which the naïve realist wishes to make, namely that the familiar properties we attribute to objects in the external world are, in the right circumstances, genuine *constituents* of conscious perceptual experience.

According to Martin (2010), to say that *o* looks *F* (in relevant cases) is to make the comparative claim that *o* has a look which is relevantly similar to a characteristic look of *F* things. It is quite consistent with this claim that *o* is not itself *F*. Identity is not the only mode of similarity. Moreover, it is quite consistent with this proposal that the look of the object in question is a property which it actually instantiates. In line with this thought, Martin proposes a distinctively parsimonious account of looks on which the looks of objects are their visually basic properties—their "size, shape, colour, visible texture, spatial arrangement of parts" (2010, p. 207)—or constructions out of these.[6]

Consider then the illusion whereby a straight stick looks bent when partially submerged in water. Applying Martin's account to this case, we might propose that it is the stick's actual colour, visible texture and straightness which constitute that aspect of its look which is (in the watery context) relevantly similar to a characteristic look of a bent stick. In itself this account is neutral concerning how we should think of a subject's perceptual experiences in seeing such a stick. The only constraint is that some visually relevant dimension of similarity obtains. Representationalists can thus accept the framework and contend that the reason that the partially submerged stick's look is relevantly similar to a characteristic look of a bent stick is that partially submerged straight sticks tend to elicit experiences in which they are *represented* as bent, which is how unsubmerged bent sticks are characteristically represented as being.

However, this is not the only explanation of the relevant dimension of similarity between the submerged straight stick and a bent stick. Martin suggests another possibility in the following passage:

> In a circumstance in which one did confront a bent stick in good lighting, a paradigmatic circumstance for encountering the shape of being bent, one would be inclined to recognize

[6] Martin also suggests that the visually basic properties include "such nonobservational properties as solidity" in a somewhat technical sense. For discussion see his 2010, pp. 206-7.

> the object as being bent; it would strike one as similar to bent things, and one would find in it an obvious similarity with bent things and a contrast with other shapes. If the psychological situation a subject is in when he or she truly utters ["The stick looks bent to me"] is relevantly similar to this paradigm kind of circumstance, then the subject is inclined to find the shape before one as similar to the paradigm of bent things—as more like being bent than anything else. (2010, pp. 214-15)

Here the dimensions of similarity which Martin highlights concern how things "strike one" in seeing the partially submerged stick, for instance the fact that one has similar recognitional inclinations to those characteristic of an encounter with a bent stick.[7]

Although the recognitional responses just mentioned are themselves naturally conceived of in representational terms, there is no commitment to explaining these responses in terms of visual experiences with representational content. Representationalists will offer this proposal but we need not accept it. As a result, by denying that the similarity underlying the relevant looks-judgments in cases of illusion commits us to shared experiential representational content, the naïve realist has an in-principle way of reconciling her approach with cases where objects look other than they are. In particular, the naïve realist can maintain that when we see a straight-stick submerged in water it is its actual look—on Martin's view, a construction from its actual visually basic properties—which is manifest to us in experience. An illusion occurs because in the relevant circumstances such a look strikes us as being similar to the characteristic look of bent things, and so *is* relevantly similar to such a look in this subjective regard. Bentness is not on this account required to explain the appearances. The actual features of the stick *in situ* suffice.[8] Martin's framework thus opens

[7] Cf. Brewer who emphasizes "a range of levels of more or less sophisticated *registration* of [visually relevant] similarities in behavioral, imagistic and conceptual categorization" including registration constituted by instinctive "reliable systematic sorting *behaviour*" (2013, p. 429, emphasis in original).

[8] Millar seems to miss this when he writes: "Martin (2010) characterizes looks as mind-independent properties of objects, but because he identifies such properties with basic visible properties such as shape, size, and colour, appealing to looks as Martin understands them would not help the naïve realist provide an account of illusion." (2015, p. 617, fn. 17)

up the possibility of explaining illusions by appeal to similarities between actually perceived features and non-actual features in a way that is quite consistent with naïve realism.[9]

Can this approach to illusions be fleshed out in ways which are empirically plausible? In the next section I critically examine Brewer's specific proposal regarding the Müller-Lyer which can naturally be located within Martin's framework. In the remainder of the paper (§§4-7), I turn to a more general approach grounded in detection theoretic models of illusion. Whilst I argue in §3 that Brewer's specific claims about the Müller-Lyer are empirically problematic, later discussion shows that his more general approach to illusions needs taking seriously from an empirical perspective.

3. A First Foray into Empirical Work on Illusions: Brewer on the Müller-Lyer

One well-known account of illusions offered by a naïve realist is that of Brewer (2008, 2011). In this section, I consider the best known of Brewer's proposals, namely his account of the Müller-Lyer illusion (Fig. 1). Brewer's proposal runs as follows:

> The [Müller-Lyer] diagram … has [visually] relevant similarities with a pair of lines, one longer and more distant than the plane of the diagram, one shorter and less distant; and those lines in themselves are a paradigm of inequality in length. In this sense the two lines look unequal in length: it is perfectly intelligible how someone seeing it might therefore take that very diagram as consisting of unequal lines, regardless of whether or not she actually does so. (2011, p. 102; also 2008, p. 176)[10]

Although Brewer's own framework is somewhat different, his view can naturally be embedded within Martin's framework. The proposal is then that the Müller-Lyer lines'

[9] Related, though importantly and variously different, approaches can be found in Travis 2004, Noë 2004, Hyman 2006, Kalderon 2011, Genone 2014 and, of course, Brewer 2011.

[10] Interestingly, in Brewer's earliest discussion of these issues (from a naïve realist, "Object View" perspective) he suggests a view closer to the confusion views discussed in §6 (see Brewer 2004, p. 70).

unequal look is a matter of their having a look (i.e. size, shape, colour and spatial configuration) which is relevantly similar to a paradigm look of inequality, viz. the paradigm look of unequal lines at different depths.

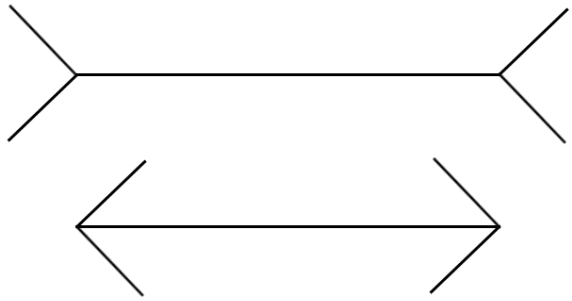

Fig. 1. The Müller-Lyer illusion (Müller-Lyer 1889). The lines are of equal length, yet notoriously the top line with outward fins looks longer than the bottom line with inward fins.

A notable feature of Brewer's proposal is that it is made in apparent abstraction from the extensive body of empirical work on the Müller-Lyer and related illusions undertaken over the last century and more. Though he does not cite him, Brewer's proposal is a version of a famous hypothesis due to Gregory (1963, 1964) based, in turn, on a more general speculation about the role of depth cues in illusions due to Tausch (1954) and Thiéry (1896). As is to be expected, Gregory's hypothesis has been subject to critical empirical scrutiny ever since it was first proposed. Given that such scrutiny has uncovered multiple grounds for dissatisfaction, it is curious that such criticism has largely been ignored by Brewer's critics.

One critic who does complain against Brewer on empirical grounds is Millar who notes that "the arrows can be replaced with circles and the illusion is unaffected" (2015, p. 619). Evidently the thought is that, since the circles do not provide depth cues, Brewer's account

of the traditional illusion is implausible.[12] However, for such a consideration to have significant force, it needs to be assessed within a broader methodological and evidential context.

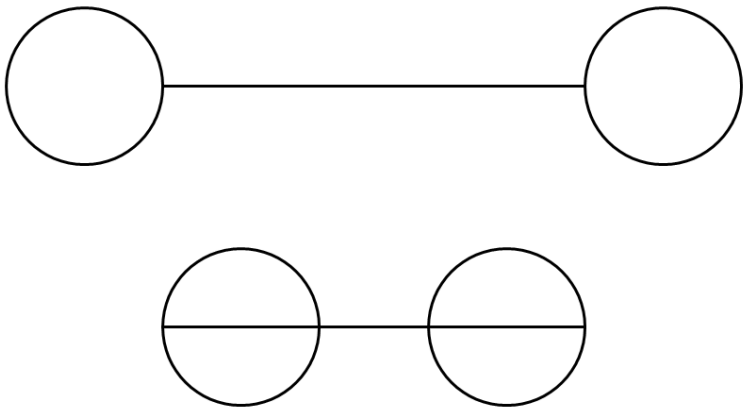

Fig. 2. The Dumbbell Illusion (Sanford 1897). Again the horizontal lines are equal in length but the top line looks longer than the bottom line.

The illusion which Millar draws attention to is known as the "dumbbell illusion" (Sanford 1897; Fig. 2) and is commonly mentioned as a concern for Gregory's account within the empirical literature.[13] However, on its own, the illusion is not usually taken to be a strong consideration against Gregory's proposal. First, in isolation, the dumbbell illusion shows at most that "Gregory's theory cannot account for all illusions" (Waite and Massaro 1970, p. 733). No doubt the illusion would constitute discriminating evidence in favour of a rival hypothesis, if that rival hypothesis were able to account for all the illusions which Gregory's account can explain and the dumbbell illusion in addition. But Millar, at least, mentions no such rival. Moreover, many theorists doubt that any general, unified account of the Müller-

[12] Antony (2011, p. 32) also points to the importance of this "version" of the Müller-Lyer, citing unpublished work by McLaughlin.

[13] An early example is Day 1965. Variants on the traditional Müller-Lyer illusion have a long history in debates about the illusion. For example, Delboeuf (1892) offers examples with triangles, squares and circles against Brentano's attempt to explain the illusion. Indeed, Müller-Lyer (1889) himself presents some fifteen variations on his illusion, including versions with curved brackets instead of fins, versions without shafts and versions without fins. All of these might, in principle, be pointed to as problematic for Gregory's hypothesis.

Lyer and all its many variants is to be had. As Woloszyn (2010) argues, citing myriad other variants of the illusion:

> [T]he existence of the illusion under such a wide variety of conditions virtually precludes the notion of a single mechanism governing all of them, given the wide range of stimuli and sensory modalities within which it appears. Researchers, therefore, might be wise to entertain the possibility that there are multiple means of producing what appears on the surface to be a single illusion, instead of continuing to pursue a Grand Unifying Theory for [the Müller-Lyer] in all its various disguises. (2010, p. 106; see likewise, Mundy 2014, p. 13 final paragraph)

If Woloszyn is right, the failure of Gregory's account to extend to all superficially similar illusions is a failure shared by all accounts. As a result, we should be open to the possibility that Gregory's account partially or wholly explains the traditional Müller-Lyer despite not accounting for superficially similar illusions.[14]

Second, insofar as the core suggestion made by Gregory and Brewer is that the Müller-Lyer arises because of misleading depth cues, it is, in fact, perfectly possible that the dumbbell illusion *can* be accounted for in a broadly similar manner. For example, we might hypothesise (cf. Woloszyn 2010, p. 103) that the dumbbell illusion arises because the line in the top figure in Fig. 2 is seen as (amodally) extending underneath the circles (i.e. as being partially occluded by them). As such it will be assigned to a more distant depth plane, and so, as with the Müller-Lyer line with outward fins, appear longer than its sibling which is arguably seen as overlaying the circles (and so as in a closer depth plane).

This all said, the evidence plausibly does weigh against Gregory's hypothesis as a major part of the explanation of the Müller-Lyer. Consider just three difficulties. First, the Holding

[14] Indeed, the situation may be even finer grained than this. Sekuler and Erlebacher (1971) argue that the effects of outward and inwards fins are a result of two quite different mechanisms.

Illusion (Fig. 3). Here the two lines appear to be offset in the plane of the paper. It is not unreasonable to expect some commonality in explanation of the Holding and Müller-Lyer illusions, yet treating the fins as depth cues cannot obviously account for the offset effect.

Fig. 3. Holding Illusion (Holding 1970). The two lines appear offset in the plane of the paper. This cannot obviously be explained by thinking of the fins as acting as misleading depth cues.

Second, DeLucia and Hochberg (1991) elicit illusions exhibiting very similar parameters to the traditional Müller-Lyer (a) using three-dimensional stimuli under free viewing conditions (experiments 2, 3 and 5), and (b) using either two- or three-dimensional stimuli without any connecting lines and/or at unnatural angles (experiments 3-5). These findings are very difficult to explain on Gregory's hypothesis. Third, and relatedly, consider the variants of the Müller-Lyer in Fig. 4. As these examples illustrate, the Müller-Lyer illusion increases in strength with extreme angles despite such angles almost never being encountered "in the wild" (and so despite not evidently affording any natural similarity with lines at depth). In contrast, as we increase the length of the fins in the fins out version of the illusion, the illusion begins to decrease once the fins are longer than half the shaft's length. Yet this is a very natural configuration to encounter. As Pressey (2013) comments: "look at the ceiling where it meets the wall in your room and focus on a corner. One is able to see the edge for a distance that much surpasses the height."[15]

[15] I quote here from Pressey 2013—a brief opinion piece vociferously expressing the author's dissatisfaction with Gregory's hypothesis. For details on extreme angles and large fins, see Pressey and Martin 1990. For further critical work on Gregory's hypothesis (in addition to that cited above, and amongst much else) see: Brown and

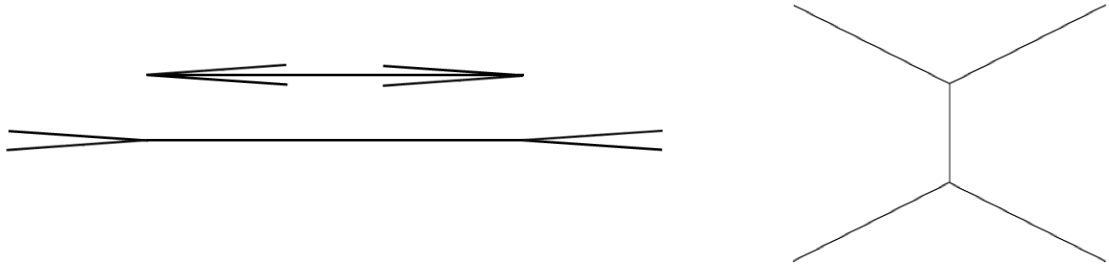

Unnatural extreme angles increase illusion **Natural large fins decrease illusion**

Fig. 4. Variants of the Müller-Lyer with large angles (left) and fin-lengths (right).

There is good reason then to be sceptical of Gregory's and so Brewer's hypothesis about the Müller-Lyer. However, this hardly spells the end for Brewer's general approach to illusions, only his specific choice of illustration. For Brewer's *general* approach to be imperilled we would have to establish that the correct account of the Müller-Lyer did not involve appeal to the idea of visually relevant similarities between the diagrammed lines and paradigms of inequality. It would be rash to think that the failure of Gregory's depth-cue account showed that.

Our initial foray into the empirical literature counsels caution in another respect. For whilst we have glimpsed some of the rich seam of work on illusions which philosophers might hope to mine, we have also glimpsed the complexity of the phenomena. Indeed, even such an extensively studied illusion as the Müller-Lyer is not fully understood. Indeed, Prinzmetal et al. confess: "We frankly have no idea of the cause of the Müller-Lyer illusion." (2001, p. 107) This lack of understanding is a consistent bugbear in the literature. In the seventies, Sekuler and Erlebacher lamented "the present pitiful state of the art of understanding illusions" (1971, p. 485). It is not clear how much things have improved. Mundy complains that "there is still no consensus of explanation [concerning the Müller-Lyer] within the literature, particularly as many theories fail to explain various modifications of the basic illusion" (2014, p. 9). Given these complexities, philosophers

Houssiadas 1965; Humphrey and Morgan 1965; Waite and Massaro 1970; Green 1972; Rock 1984; Day 1989; Morgan et al. 1990; Howe and Purves 2005; and Woloszyn 2010.

might be forgiven for beating a hasty retreat to the armchair. Is there a way of pressing on in such circumstances?

Confronted with complex phenomena, a familiar strategy is to build simple models. As Williamson (forthcoming) describes:

> When a system resists direct study, because it is so complex or hard to observe, model-building constitutes a key fall-back strategy. Studying a model often yields insight into the phenomena it models.… macroscopic phenomena are typically too complex and messy to obey many informative exceptionless generalizations framed in macroscopic terms.… it may be more realistic and more fruitful to aim at building increasingly good models instead.

Illusions are highly complex phenomena. Unsurprisingly, then, scientists have sought to build, study and test simple models in order to understand them better. Amongst the most important psychophysical models of perception developed in the last century are those of signal detection theory (SDT). Such models offer an opportunity to consider philosophical disputes about illusions without commitment to a detailed understanding of the mechanisms generating particular illusions. I pursue this approach in what follows. First, I sketch the basic SDT model of a simple discrimination task, emphasising how it distinguishes two aspects of a perceiver's responding: their discriminative sensitivity, and their criterion or bias. I then turn to actual applications of this model to geometric illusions. Drawing on work by Morgan et al. 1990 and Witt et al. 2015, I explain why the upshot of this application might at first sight seem to tell in favour of representationalist accounts of illusions (§5). Then, in §§6-7, I explain how further reflection reveals at least two alternative interpretations of the model, both in different ways congenial to naïve realism.

4. A Simple SDT Model of a Sensory Discrimination Task

In this section I sketch a SDT model of a simple sensory discrimination task. It is important to emphasise that I am not putting forward this model as a true, let alone complete, account of perceptual discrimination. Indeed, not least from a naïve realist perspective, there are clear problems with certain assumptions embedded in the model.[16] It is better to think of the model as akin to models of mechanical and electromagnetic systems in physics which falsely assume that masses and charges are possessed by "point particles". Much can be learnt from such simplified models. In a similar way, the present aspiration is to learn something about illusions by studying simple models even if these models involve false simplifying assumptions. In particular, such models help bring out certain key distinctions (for example, as I explain below, between shifts of perceived feature and shifts of perceptual response, or as between natural and unnatural response criteria) which I argue are of central importance in thinking about illusions.

The simple discrimination task to be modelled, whose relevance will shortly become apparent, involves just two stimuli: a "short" 5cm line and a "long" 7cm line. On each trial one line is presented, and the observer must classify it as either "short" or "long". SDT models this task by associating each (type of) stimulus with a distribution of sensory effects. In the simplest such models these distributions are assumed both to be normal and to have equal variance. Clearly, to the extent that an observer is sensitive to differences in length, these distributions will be different for lines of different lengths. If 5cm and 7cm lines elicit the very same pattern of effects on a given system, that system will be entirely *in*sensitive to their difference in length. A natural measure of the *discriminative sensitivity* of an observer in the present task is the distance between the means of the 5cm and 7cm distributions.

[16] In particular, the model takes for granted the idea of stimuli eliciting "sensory effects" along a single dimension, and moreover that different types of stimuli can, with varying probabilities, elicit one and the same type of effect. On the face of it this is to embrace a "common factor" approach to perception which the naïve realist will reject. This is perhaps not surprising given that the original problem of "signal detectability" which motivated SDT was that of characterizing and optimizing the method which a radar operator "given a voltage varying with time during a prescribed observation interval" should use "to decide whether its source is noise or is signal plus noise" (Peterson et al. 1954, p. 171). The large question of how to (re)interpret or amend our simple model of perceptual discrimination in such a way that is consistent with naïve realism lies beyond the scope of the present paper.

Knowing how sensitive an observer is does not suffice to determine their performance on the task in question, however. This is because, in responding, subjects must implicitly operate with a criterion or standard by which to judge when a line should be classified as "long" (and correspondingly "short"). In the present case, it would seem natural for subjects to adopt a criterion roughly midway between the two distributions' means, i.e. a criterion corresponding to the mean of the hypothetical distribution of sensory effects associated with a line roughly 6cm in length (Fig. 5).

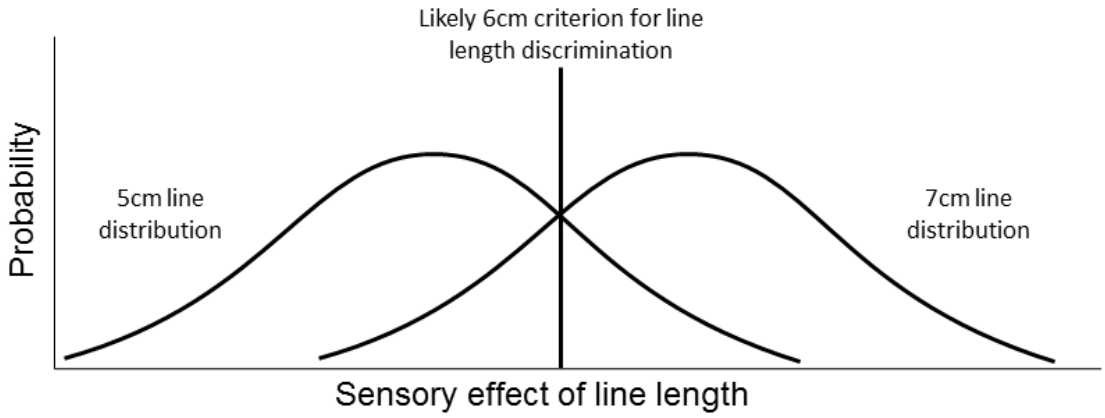

Fig. 5. An idealised SDT model of a simple discrimination task showing a plausible response criterion which a naïve observer might exhibit.

In many tasks we will want to consider multiple stimulus-types and multiple categories of response but this does not greatly complicate the analysis so long as discrimination is all along a single dimension. Where there are multiple stimuli, we simply associate each stimulus type with its own distribution, again distinguishing between subjects' sensitivities (given by distances between the means of these distributions) and response criteria (given by a single threshold in the case where all stimuli are being categorized into two classes, and multiple thresholds where there are multiple classes).[17] The crucial point to emphasise is that

[17] For further details see Macmillan and Creelman 2005, chpt. 5. The classic text on detection theory is Green and Swets 1966. The appendix to Palmer 1999 provides a brief but helpful introduction.

in all these cases performance is modelled in terms of two parameters—sensitivity and criterion, commonly also known as bias.

To illustrate the effects of altering these parameters in our simple task, consider first a subject whose criterion shifts to the right in Fig. 5. This means that they will show a decreased tendency to categorise lines as "long". Whether this is a bad thing depends (a) on the prior probabilities of the line being 5 or 7cm and (b) on what matters to the observer. One thing that might matter to the observer is maximising the number of occasions on which they correctly classify a line (their percent correct). Then, assuming, that the two types of line are presented equally often, the optimal choice of criterion will be midway between the two distributions (as in Fig. 5). If, however, most presented lines are 5cm long, then a rightward shifted criterion will be optimal. Furthermore, it may be significantly more important to an observer to correctly classify short lines as opposed to long lines. (Perhaps they earn 50¢ for correctly classifying a short line but only 1¢ for correctly classifying a long line.) Then, even if the lines are equally probable, the optimal (payoff maximizing) strategy will be to operate with a rightward bias.[18]

It is commonly thought that a major boon of signal detection theory is that it allows us to separate out these two otherwise confounded aspects of an observer's responding: their sensitivity and criterion. As Green and Swets write:

> [A] principal advantage of modern detection theory is that it shows how to compress a host of factors which affect the observer's attitude into a single variable, called the decision or response *criterion*, and how to use false-alarm responses to estimate the level of the criterion.

[18] To take a more familiar example, imagine that you are a bouncer at an over-21s night club. Your job is to spot underage clients and ID them. Imagine 50% of your queue are under-age. If your aim is to maximize correct identifications of underage clients ("hits") whilst avoiding more than a small percentage of "false alarms" (incorrect identifications of over-21 clients as underage), that means you will only ask for ID from those who look significantly under age (i.e. you will be *conservative* in asking for ID). If you do this, you will still make some mistakes but most of those whom you ID will be underage. If instead your aim is to avoid "misses" (failures to identify underage clients as such), then you will adopt a very strict policy, ID-ing anyone who might conceivably be underage (i.e. you will be *liberal* in asking for ID).

> By extracting two parameters from the data—one related to attitude and one to sensitivity … the procedures of detection theory isolate nonsensory factors, so that a relatively pure measure of sensitivity remains. (1966, pp. 118-9)

Likewise, Macmillan and Creelman describe how Green and Swets "prescribed experimental methods and data analyses for separating decision factors from sensory ones" (2005: xiii). All this strongly suggests that SDT is conceived as offering the tools for distinguishing between sensory, experiential contributions to responding and decisional, cognitive contributions. On this conception, criteria setting/bias is a purely decisional, non-perceptual matter, whereas perception is exclusively characterized in terms of discriminative sensitivity.

5. Illusions, Biases and an Apparent Vindication of Representationalism

The SDT model just sketched not only allows us to characterize sensitivity independent of bias but also to investigate whether particular interventions affect sensitivity or bias or both. In keeping with the assumption just flagged at the end of §4, viz. that bias is decisional and sensitivity perceptual, a standard application of the model is to investigate whether a given effect is perceptual or cognitive.[19] This application has a striking implication in relation to paradigm cases of illusion. In particular, studies of the Müller-Lyer (Morgan et al. 1990, Witt et al. 2015) show that the addition of inward or outward fins to straight-lines exclusively shifts subjects' criteria without affecting their sensitivity. Following the standard interpretation of bias and sensitivity above, this implies that the Müller-Lyer is not a perceptual effect!

To see how this is a *theoretical* possibility, return to our simple line length discrimination task, but now add outward and inward fins to the 5cm and 7cm lines. In such an experiment subjects show an increased tendency to classify lines with outwards fins as "long" and a

[19] See the various examples highlighted in Witt et al. 2015 such as Grove et al. 2012 on the sound-induced visual bounce effect (Sekuler et al. 1997).

decreased tendency to classify lines with inwards fins as "long". Sensitivity remains unchanged. The apparent upshot is that the addition of outward or inward fins shifts subjects' criteria: inward fins induce a bias towards responding "short"; outward fins towards responding "long". This is illustrated in the top row of Fig. 6 which shows how shifts in response criteria corresponding to "fins in" and "fins out" conditions respectively can explain the differing patterns of judgments made by subjects in the discrimination task.

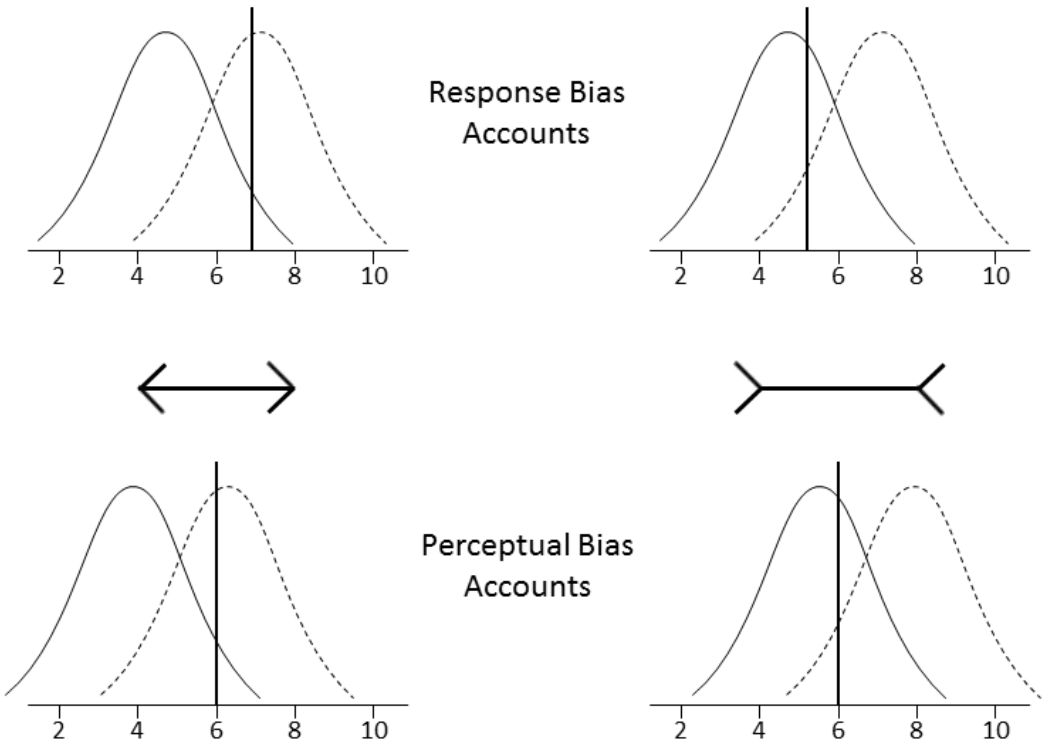

Fig. 6. Figures in the top row illustrate how the Müller-Lyer illusion can be explained in terms of shifted response criteria for the "fins in" (LHS) and "fins out" (RHS) versions. Figures in the bottom row illustrate how the Müller-Lyer illusion can be explained in terms of perceptual bias, i.e. in term of shifts in the distributions of sensory effects. This is shown for the "fins in" version (LHS) and "fins out" version (RHS) respectively. Notice how the criterion remains unchanged at the hypothetical natural criterion of 6cm in both of the bottom figures. Notice also how sensitivity is constant across all cases. Figure based on Witt et al. (2015, p. 294).

Despite the capacity of this purely cognitive account to model the data effectively, the conclusion that the Müller-Lyer is not a genuinely perceptual effect is regarded by both Morgan et al. (2012) and by Witt et al. (2015) as a *reductio* of the way in which SDT has been interpreted. It is a *reductio* because the Müller-Lyer is "an undoubtedly perceptual effect" (Morgan et al. 2012, p. 186). Neither group abandons detection theory, however. Rather both argue that the notion of "bias" should never have been exclusively associated with judgment and decision-making. Bias they argue also encompasses *perceptual* bias. Paradigmatic illusions such the Müller-Lyer which do not involve any change in sensitivity can nonetheless be perceptual insofar as they involve such *perceptual* bias.

To see how perceptual bias is to be understood consider the bottom row of Fig. 6. This shows how we can explain the Müller-Lyer illusion without altering a subject's response criterion *or* their sensitivity by thinking of the fins as shifting both of the distributions associated with the two line-lengths. Inward fins induce a leftwards shift which lowers the probability of a "long" response; outward fins induce a rightward shift which increases the probability of a "long" response. As can be seen from casual inspection, such an approach succeeds in capturing the data just as successfully as an approach which appeals to response bias. However, the understanding of the illusion on a perceptual bias account is very different. As Witt et al. write, "when the tails are oriented inwards, this creates a perceptual shift to see both short and long lines as shorter that they would otherwise be perceived" (p. 295).[20] In contrast, on the earlier response bias model, in respect of their length-related sensory effects the lines with fins are *perceived* just as lines without fins are (witness the sameness of sensory distributions associated with each). They are nonetheless *judged* differently because of the shifted decision criteria which the fins induce.

This brings us finally back to the question whether empirically-based models of illusions have implications for the dispute between naïve realism and representationalism. In light of

[20] The same understanding is suggested by Morgan et al. (1990). Witness their opening question: "Do the fins of the Müller-Lyer illusion *change the perceived length of the line only*, or do they in addition decrease the observer's sensitivity to length differences?" (1990, p. 1794; see also p. 1795, again my emphasis)

the discussion hitherto, it is tempting to think that they do. Empirically adequate SDT models of geometric illusions such as the Müller-Lyer appear subject to two interpretations. On the first interpretation they result from pure shifts of *decision* criteria. This interpretation is consistent with naïve realism yet implausibly conceives of illusions as purely cognitive, non-perceptual effects. The second interpretation appeals to the idea of perceptual bias. This, as the name suggests, avoids treating illusions as purely cognitive phenomena. However, the interpretation is liable to seem inconsistent with naïve realism insofar as perceptual bias is understood in terms of our perceptual system shifting its response to one feature (e.g. line-length) such that it responds in a manner appropriate to a different feature (e.g. a longer or shorter line-length). (See again the bottom row of Fig. 6.) This accords naturally with a representationalist picture on which illusions are understood in terms of our perceptual system responding to an object with a given feature in a manner appropriate to a different feature, namely by *misrepresenting* the object as having a different feature. In contrast, it is not obvious how to make sense of perceptual bias on a naïve realist account which denies that illusions involve systematic perceptual *misrepresentation*. In short then, our initial examination of a SDT model of the Müller-Lyer illusion seems to tell in favour of a representationalist account of the illusion and against a naïve realist approach.

In what remains of this paper, I challenge this overhasty conclusion by offering two alternative ways of thinking about illusions on the model before us. These two responses respectively challenge the interpretations of perceptual and response bias adopted above. The approaches are neither exclusive nor exhaustive. They are also *modest* in the sense that they operate within the basic modelling strategy which SDT offers, challenging instead the model's interpretation. More radical approaches are eminently possible.

6. In Defence of Naïve Realism I: Perceptual Bias Reconsidered

Does understanding illusions in terms of perceptual bias really commit us to representationalism? Consider so called "confusion" (Woodworth 1938, p. 645; Erlebacher and Sekuler 1969) or "incorrect-comparison" (Rock 1984, p. 167) approaches to the Müller-

Lyer. On this family of views, the illusion arises, as DeLucia and Hochberg put it, because observers respond, at least in part, "not to the distance they are asked to judge, but to other dimensions of the figure instead" (1991, p. 553). Such theories are closely related to Pressey's "assimilation" theory (Pressey 1967, 1971) according to which the line length judgements in the Müller-Lyer are always made in the context of other magnitude judgements and are biased towards the mean of these other judgments. They are also related to, though importantly distinct from, so-called "perceptual compromise" or "conflicting cues" theories (Day 1989). On these theories the line's length is *misrepresented* because our perceptual systems draw on both local cues to line length but also conflicting information concerning other aspects of the figure, effecting a compromise between these two conflicting sources of information in representing the line's length. Though the differences between these theories is important (if not always fully explicit), all emphasise the relevance of other dimensions of the Müller-Lyer figure in addition to the length of the connecting shaft whose length is, of course, the intended target of subjects' judgments.

Different theories emphasize different "other" dimensions. Some confusion theorists point to the distance between the tips of the fins as opposed to the ends of the central line (Erlebacher and Sekuler 1969). Others point to the distance between the (mean) geometrical centres of the arrow-heads (Morgan et al. 1990, Morgan and Glennerster 1991). Here I focus on a suggestion most prominently associated with Day (1989), who as just noted is a "compromise" not a "confusion" theorist. I intend this discussion to be illustrative of a possible "confusion" position and without commitment to the inevitably complex details.

According to the suggestion to be explored, a crucial clue to the Müller-Lyer illusion is the fact that the visual system is interested both in *global* information about objects as wholes and their overall features, as well as more *local* information about the parts of objects and their detailed features. As Navon (1977) discusses in his classic paper, the visual system is typically biased towards global as opposed to local features. To show this Navon created special figures: large letters comprised of smaller letters (Fig. 7). In one experiment, subjects

had to make speeded identifications of either the small letter elements or of the large overall letter. Navon found that subjects who had to identify the small letter elements were slowed when the large letter was inconsistent. In contrast, subjects who had to identify the large letter were not slowed when the small letter elements were inconsistent. This, he concluded, shows that "whereas people can voluntarily attend to the global pattern without being affected by the local features, they are not able to process the local features without being aware of the whole" (ibid., p. 371).

Fig. 7. Examples of consistent and inconsistent "Navon" letters. From Jansari et al. 2015. Reproduced under Creative Commons Attribution License (CC BY).

This "global precedence" bias has subsequently been explored in a wide variety of contexts. One way to make sense of it is to think of the visual system's adaptive purpose as being to locate, assess and categorize biologically relevant objects as such, be they tools or trees, predators or plants, boulders or hiding holes. In visually tracking such objects it is arguably their global features which matter first and foremost. It matters to identify rapidly how large a potential predator is, whether a hole is big enough to hide in, or how sizeable a meal a plant will afford. This is of course not to say that local features are irrelevant, only that the visual system will naturally prioritize global information in parsing the scene.

Might this help us understand the Müller-Lyer? Adopting the above perspective, we might note that the two Müller-Lyer figures, whilst matched in local line length, plainly differ in

global size.[21] Fig. 8 offers a crude visualization, though it is important to recognize that this is purely illustrative: how the visual system tracks "global" size and shape is a complex and controversial issue.[22]

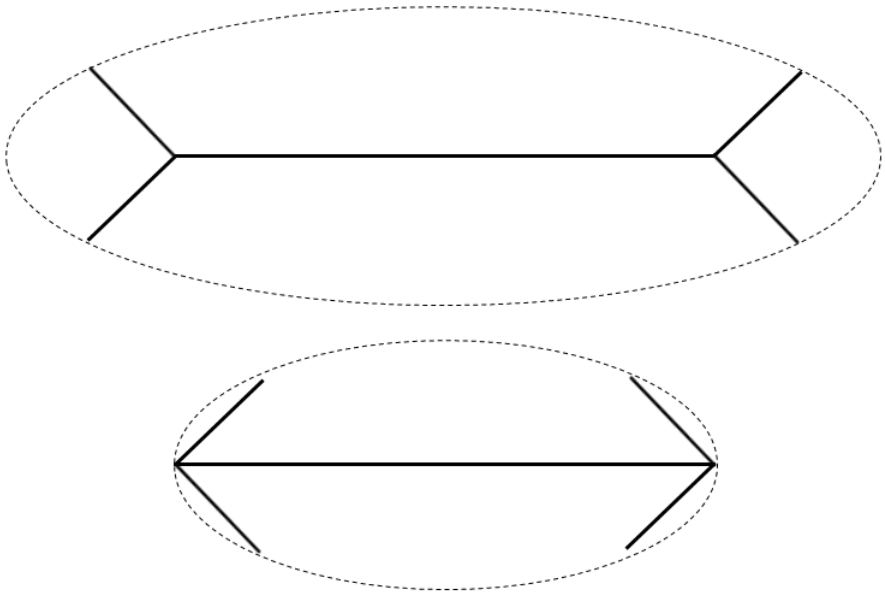

Fig. 8. Illustration of the suggestion that the visual system might track the two versions of the Müller-Lyer as having different "global" sizes. Figure based on Mundy 2014, p. 10, panel B. Mundy is discussing Day 1989.

If the visual system is biased towards tracking the lengths/sizes of objects as wholes, then we might think of perceptual biases not as shifting the perceived lengths *of the lines* but as shifting *which features* are underpinning subjects' responses. Slightly more accurately, we might think of a simple line without fins as a special case where global size coincides with local line length. The addition of fins thus introduces a dissociation between the (global)

[21] This is also true of the dumbbell illusion as well as other, e.g. three-dimensional, variants of the Müller-Lyer. As a result, the current proposal appears to avoid many of the criticisms ranged against Gregory's in §3.

[22] Indeed, a moment's thought reveals that Fig. 8 does not suffice to explain the difference in length judgments between a simple line and the line with inward fins. However, as Rock (1984, p. 167) notes the preponderance of the traditional comparative illusion actually comes from the fins out version. Furthermore, as I noted above, there are various versions of the confusion hypothesis. On one, the confusion is between the distances between the (mean) geometrical centres of the arrow-heads as opposed to the distances between the line-termini independently of the arrow-heads. This global bias explains both versions of the illusion. Here see Morgan et al. 1990 and Morgan and Glennerster 1991; also Pressey 1971.

feature which is most salient and guides responding, and the (local) feature which is the intended target of those responses. The crucial point is that, insofar as subjects' responses are guided by their perceptions of the figures as wholes, the sensory distributions relevant to that judgement will be different to those associated with simple line length—and in precisely the manner of the perceptual biases discussed above.

The reader may well wish to object: surely we can see that the figures have outward or inward fins and are not confused between the global length of the figures and the local lengths of the lines! However, arguably the illusion simply shows how hard it is exclusively to base our perceptual judgements on local perceptual analysis, and abstract away from global features in making local line-length determinations. As Rock writes: "Despite a clear understanding of what parts of the line are to be compared, we cannot avoid including other components in our judgments" (1984, p. 167).[23] Nonetheless, the objection suggests a prediction: insofar as we can manipulate subjects to focus on global versus local features, the strength of the Müller-Lyer illusion should vary. There is good evidence for this. For instance, Bates (1923) gave subjects two different instructions, one encouraging them to "pay attention to the total impression", the other to adopt a "critical" or "analytical" attitude and to "pay attention to the two horizontal linear extents, abstracting as far as possible from the 'wings'" (1923, p. 65). Bates found that the adoption of this latter "analytical" attitude decreases the illusion.[24] Relatedly, Gardner and Long (1961) demonstrate that attending to the shaft whilst ignoring the fins reduces the magnitude of the illusion. Finally, Mundy (2014) manipulated subjects into focusing on global or local features by instructing them to read out either large "Navon" letters or their smaller letter elements for five minutes. After this treatment, subjects had to adjust a simple straight-line to match a Müller-Lyer line. Mundy found that "the strength of the Müller-Lyer illusion was significantly increased for participants in the global processing [large letter] bias group, in comparison to those in the control condition; and it was significantly decreased for participants in the local processing

[23] Here compare the way in which we cannot avoid having our attention drawn to, and so being affected by, the large, "global" letters in Navon's task.
[24] See further Day 1962 and Eaglen and Kirkwood 1970.

[small letter element] bias group, in comparison to those in the control condition" (ibid., 12).[25]

Discussing a related account, Morgan et al. (1990) draw the following conclusion:

> The term "illusion" is a value judgement put upon [the subject's criterion] by the experimenter, because the subject has not made exactly the judgement that the experiment intended. We have suggested that the reason for this is that the visual system is highly constrained in the nature of the judgements which it is able to make. It is possible to formulate verbal instruction which the visual system is not able to carry out exactly. The geometrical illusions are perhaps best understood, not as mistakes by the visual system, but as a failure of the visual system to carry out the exact measurement required of it [i.e. a judgment of line-length *per se*]. Our proposal is that there are … dramatic constraints on spatial vision, and that when we understand them better, the term "illusion" will no longer be necessary. (1990, p. 1809)

Here it is tempting for the naïve realist to understand the final rhetorical flourish in terms of the idea that (at least certain) illusions are not after all cases where our visual experience misrepresents objects, but rather cases where actually instantiated features of the environment are perceived. The temptation towards thinking in terms of misrepresentation arises because we make a false assumption concerning the features which account for a subject's judgements. In the case in point we assume that judgements are grounded in the perception of simple local features such as the length of the central shaft in the traditional Müller-Lyer. This assumption is mistaken, however. The features most salient to subjects and to which we must appeal in explaining their judgments turn out to be significantly global. Yet these global features are actual features of the displays in question. Thus, the naïve realist has no difficulty in appealing to them in articulating the conscious character of subjects' perceptual experience.

[25] Another related finding is that the illusion decreases when the colour and/or luminance of the shafts and fins differs (Bates 1923, Mukerji 1957, Sadza and de Weert 1984).

7. In Defence of Naïve Realism II: Response Bias Reconsidered

I now consider a second possible way of interpreting detection theoretic models which is congenial to a naïve realist. This option does not appeal to perceptual bias but instead involves accepting the idea that illusions such as the Müller-Lyer involve shifts of response criteria. This interpretation was swiftly rejected above on the ground that it involved treating the illusion as non-perceptual. Here I question this verdict.

The appeal to response bias to understand the Müller-Lyer is closely related to Brewer's approach.[26] Consider Brewer's first discussion of how illusions are to be construed on his Object View:

> "Illusions", then, are absolutely not cases in which there is some kind of misrepresentation of reality by perceptual experience. For the subjective qualities of perceptual experiences are constituted by the various features of mind-independent things that are accessible to the subject, given the relevant conditions of perception…. Rather, they are cases in which the way that the subject is most naturally inclined to judge the world to be, given which features of mind-independent reality are accessible to him in experience in this way, is systematically out of line with the way things actually are out there. (2004, p. 74)

It is tempting to read Brewer here as adopting the view that illusions are fundamentally to be understood at the level of judgement and so as non-experiential phenomena (cf. Phillips 2005). However, Brewer has since insisted that illusions really are "*experiential*" phenomena, "a matter of the phenomenology of perceptual experience *itself*" (2011, p. 119; his emphasis). In this regard, several critics have remained unconvinced (e.g. Siegel 2011, p. 66, fn. 40, and Smith 2010, p. 399, fn. 20). Their objections can be thought of as analogous to claims noted

[26] In contrast, the account of §6 is rather closer to Brewer's approach to the bent stick illusion at least in his early discussions (see Brewer 2004, p. 73; cf. Brewer 2011, p. 106).

above in Morgan et al. (1990) and Witt et al. (2015) that illusions, being genuinely perceptual effects, cannot arise purely from shifts in response criteria.

To probe this assumption, consider a closely related psychophysical study due to Morgan et al. (2012) designed to make a complementary point to that of Morgan et al. (1990). In this study, Morgan et al. aim to demonstrate that pure shifts of response criteria can occur for wholly non-perceptual reasons. They show this by demonstrating that subjects can voluntarily shift their response criteria so as to adopt an "unnatural" criterion without any significant impact on their perceptual sensitivity. Given their earlier results in relation to the Müller-Lyer, Morgan et al. argue that this shows that response shifts in the absence of any change in sensitivity are entirely consistent with both deliberate response strategies (e.g. deliberately setting out to adopt an "unnatural" criterion) but also with genuinely perceptual effects (as in the Müller-Lyer). Simply knowing that an effect is on bias and not sensitivity is quite consistent with either kind of influence.

Morgan et al. (2012) use a vernier acuity task in which three dots are presented on a given trial, with the middle dot offset varying degrees from the vertical (Fig. 9(A) and (B)). The subject must categorize this offset as leftwards or rightwards. Although Morgan et al. use a range of different offsets we can model the task in much the way we modelled the two line categorization task introduced in §4. As with all such tasks, even though no explicit standard is shown, subjects must adopt an implicit standard.

A **B** **C** **D**

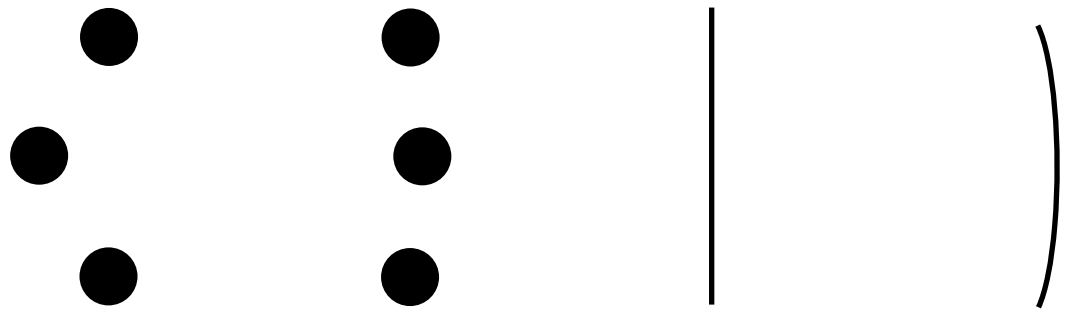

Fig. 9. Two example stimuli from a vernier acuity task are shown in 9(A) and 9(B). 9(C) represents a "natural", vertical criterion by which to judge offset. On such a criterion vertically aligned dots will be judged offset rightwards 50% of the time. 9(D) represents an "unnatural", curved line criterion. Adopting such a criterion, subjects will instead judge the dots in 9(B) as offset rightwards 50% of the time.

A crucial notion which Morgan et al. (2012) exploit in their presentation is that of a "natural criterion". There is, they suggest, "an obvious sense in which [categorization] tasks such as vernier acuity … have natural null points. Natural criteria are those that the observer can be verbally instructed to adopt without the need to show them the null point, and without the need for response feedback." (2012, p. 186) In the vernier acuity task, subjects naturally adopt a criterion corresponding to the point at which the dots are at zero offset from the vertical. In other words, their "implicit standard is a straight line" (Morgan et al. 1990, p. 1794; Fig. 9(C)). This is hardly surprising. What is it natural to compare three dots with when asked about the offset of the middle dot? Surely a straight line, not a curved line. Similarly, in the two line categorization task introduced in §4, the natural null point corresponds roughly to a 6cm line, i.e. lies midway between the distributions associated with the 5cm and 7cm lines. Again, this is not surprising. It would be strange if subjects chose a point strongly skewed towards 5cm or 7cm. That said, as Morgan et al. note, in principle, vernier "stimuli could be compared with an imaginary curved line, rather than an imaginary straight line" (2012, p. 186). Adopting this "unnatural criterion" would correspond to taking the judgment to be a matter of judging offset relative to the curved-line shown in Fig. 9(D). Notice how, judged by such a criterion, the dots in Fig. 9(B) would then be judged as rightwards offset 50% of the time.

Morgan et al. (2012) exploit the notion of a natural criterion to "define a perceptual shift as a translation of the psychometric function [i.e. the distribution of responses], which occurs without a change in the observer's natural criterion" (p. 186). As I discuss shortly, this definition is not ideally expressed. Nonetheless, it is not hard to see what Morgan et al. have in mind. Consider first their finding that subjects are able to change criteria deliberately without affecting their sensitivity to changes in dot alignment. Here Morgan et al. clearly treat the shift in responding to be a *non-perceptual* shift in criterion. Going by their definition this must mean that there *has* been "a change in the observer's natural criterion". And clearly what this must mean is that there has been a change in the observer's criterion *from a natural one to an unnatural one*. After all, the new criterion is in a clear sense "unnatural". It is either the result of the deliberate and explicit adoption of a response policy, or the result of attempting to conform to deliberately biasing feedback. It is not a criterion a subject would adopt without feedback or explicit instruction. But notice that, although there has been a change in criterion from natural to unnatural, there is no good reason to think that there has been any change in the subject's *natural* criterion itself. It remains the case that subjects will default to a straight-line criterion as soon as they drop their voluntary decision to respond in a deliberately biased manner. Indeed, since the stimuli are just the same, unless we think that adopting a policy alters the way things appear, it is very implausible to suppose that we have here a perceptual shift. The upshot is that it would be better to define a *perceptual* shift as one which occurs *without the subject adopting an unnatural criterion* and conversely a non-perceptual shift as one in which the subject *does adopt an unnatural criterion*. In this way, we tie perceptual and non-perceptual shifts to the notion of perceptually natural criteria.

Adopting this new understanding, consider again the Müller-Lyer. This is by wide-agreement a perceptual effect and moreover one which does not operate on sensitivity. As a result, the shift must be one which occurs *without the subject adopting an unnatural criterion*. However, this might occur in two quite different ways. It could occur because there is no change in criterion at all, the effect instead operating by shifting the distributions of sensory responses. This is what is proposed by Witt et al. 2015 above in apparent conformity to a

representationalist model of illusion. However, it could also occur because which criterion it is natural for subjects to adopt changes with the addition of fins.[27] On this understanding subjects do adopt a *different* criterion but they do not adopt an *unnatural* criterion. The shift of criterion can thus be considered perceptual in that it does not involve the adoption of an unnatural criterion, despite not involving any shift in the distributions of sensory effects associated with the stimuli.

This picture fits nicely with Brewer's general approach to illusions. On the one hand, the addition of fins does not lead to any change in subjects' sensory distributions. This fits with the naïve realist claim that the addition of fins does not prevent subjects perceiving the lengths which the lines actually have. Instead, the addition of fins alters the criterion which subjects adopt, and this is what accounts for their responses in the discrimination task (as in Fig. 6, top row). Crucially, however, this shift of criterion is not to be construed as a deliberate shift in responding located purely at the level of judgment. It is a *natural* as opposed to *unnatural* shift. In this way, the effect is genuinely perceptual.

The above discussion shows the need for a notion of perceptual naturalness within psychophysical work on illusions. As such it should embolden the naïve realist to rely on such a notion. Indeed, I suggest that exactly this notion is in play when Brewer argues that in being consciously acquainted with the Müller-Lyer lines (and the very lengths they actually have), a certain visually relevant similarity "jump[s] out at me or *capture[s]* my attention". The similarities that leap out are the perceptually natural ones in the relevant context. Perceptual naturalness equally provides a way of fleshing out Martin's appeal to the similarities which "strike one" as "obvious" in a given perceptual circumstance, and

[27] And not just the addition of fins. The evidence suggests an extremely complex interaction between natural criteria and many other general and subject-specific factors. Witness the following passage from a recent paper investigating the effects of orientation on the Müller-Lyer. "My experiments … with Müller-Lyer patterns were frustrating. … The results … were erratic. They were strongly subject-dependent, there was no simplifying symmetry when the patterns were turned upside down, etc. My provisional, not too satisfactory, explanation is that a subject may compare the lengths of the segments between the fins according to various criteria, (for instance, forming a virtual rectangle with a pair of segments, looking at orientations, etc.) and the criterion he/she chooses depends upon the orientation of the stimulus." (Ninio 2014, p. 14)

likewise his talk of the paradigms which subjects are inclined to find the stimuli before them as most similar to (2010, pp. 214-15).[28]

Critics of naïve realism will likely raise two concerns at this juncture. First, they will press for a clearer picture of what naturalness amounts to. Second, they will press for further justification of the claim that shifts of natural criteria constitute genuinely perceptual effects. Let me take these concerns in turn. First, can more be said about naturalness? I introduced the notion of naturalness above in terms of those criteria which "the observer can be verbally instructed to adopt without the need to show them the null point, and without the need for response feedback" (Morgan et al. 2012, p. 186). However, I suggest that this gloss is best read as a generic claim and not as a definition. Instead, the notion of naturalness should be understood as primitive in respect of personal level psychology.

Can any more be said? Need it be? It is not at all clear that the naïve realist should accept the demand that they explain at a personal level why certain stimuli strike subjects in the ways that they do, or why certain similarities are natural as opposed to others. Such explanations may only be available at lower levels of explanation—by thinking about the processing and design of the visual system. If this is right, the naïve realist is entitled simply to say that certain similarities are perceptually natural for subjects, pointing to vision science for further explanation as required. Representationalists cannot object to this gambit for they offer exactly the same response in relation to an analogous question, namely: why in cases of illusion are the perceived objects misrepresented as being *F* when they are not-*F*? Representationalists will surely want to direct us to vision science for answers. They certainly need not accept the demand that they provide an explanation at a personal level of explanation. It is unclear then why it should be any different with respect to the question: why in cases of illusion do the perceived objects *strike* subjects as being subjectively similar to *F* objects despite themselves not being *F*?

[28] It is interesting to compare Craig 1976 on the notion of "not being able to help thinking of something as being such-and-such" (p. 18, and §IV *passim*).

Second, can anything more be said to justify the claim that shifts of natural criteria constitute genuinely perceptual effects? Behind this question lies another problematic assumption. This is the assumption that, unless there is some difference in which properties are perceived or represented, there can be no perceptual difference.[29] Contemporary naïve realists reject this assumption, insisting that perception is a three-place relation between subjects, objects and "perspectives" or "standpoints" (e.g. Campbell 2009, Brewer 2013, and French Forthcoming). This third ("standpoint" or "perspective") relatum is intended to capture the fact that we can consciously perceive features in a host of different *ways* depending on the circumstances and idiosyncrasies of our perceptual situation. To be perceptual, then, an effect need not involve a difference in *what* is seen, it might instead involve a difference in *the way* in which features are seen. Given this, the naïve realist can reasonably propose that a case in which a subject adopts a different *natural* criterion is a case in which they perceive a given feature in a different *way*. In contrast, all else being equal, a case in which a subject adopts an *unnatural* criterion, is not a case in which there is any difference in the way the subject perceives the feature, but only in the way in which they post-perceptually respond to it. The upshot is that there is no good reason to deny that shifts of natural criteria constitute genuinely perceptual effects.

8. Conclusion

Despite initial appearances, empirically grounded models of illusions do not favour the representationalist. Indeed, thinking about such models and associated psychophysical

[29] This assumption (sometimes known as "diaphaneity") harks back to the early twentieth century sense-datum theorists. Thus, Price (1932, p. 5): "Are there several different sorts of acquaintance…? I cannot see that there are. The difference seems to be wholly on the side of the data." The assumption is arguably also embedded in the traditional gloss on SDT according to which bias is a purely decisional parameter. Behind that idea is the thought that there is a unique perceptual experiential state corresponding to any given value of the sensory effect parameter, and that all other variation is non-sensory or decisional. By interpreting the location of a subject's natural criterion as an aspect of their perceptual situation I am in effect rejecting this interpretation (though in a rather different way to Morgan et al. and Witt et al. above). It is interesting to consider the extent to which the pioneers of perceptual SDT are influenced by sense-datum theorists. Note how Swets et al. clearly echo the language and presumptions of Moore (1953) and Russell (1912) when write that they shall "use the term *observation* to refer to the sensory datum on which the decision is based" and "assume that this observation may be represented as varying continuously along a single dimension" (1961, p. 304).

work reveals two ways of understanding illusions which are quite congenial to the naïve realist. On the first, illusions arise because of the salience impact of features other than those about which we are directed to make (and take ourselves to be making) judgements. On the second, illusions arise because of the effects of contextual cues on which similarities are perceptually natural. Which (if either) of these accounts applies to which (if any) illusions is a matter for future (and largely empirical) investigation. The moral here is that at least with respect to certain familiar illusions (the Müller-Lyer and its variants but also, I suggest, other well-known examples such as the Delboeuf, Ponzo and Jastrow illusions) the naïve realist need not deny that the relevant perceptions are of the same basic kind as ordinary veridical perception (*pace* Foster and others).

Of course, it is a large and further question whether these approaches can be extended to the wide variety of other phenomena classified as illusions. Given their variety, we should not expect any "one size fits all" approach. And, as already acknowledged, some cases plausibly do require a disjunctivist account. Nonetheless, the present investigation casts doubt on the representationalist claim to have an eminently better account of illusions.[30] It also undermines the aspersion that naïve realism is an antediluvian view inconsistent with basic science. On the contrary, for the naïve realist, there is real value in engaging with the science of illusions. For such work reveals that illusions may not be all that they first seem.[31]

References

Antony, L. 2011. The Openness of Illusions. *Philosophical Issues* 21(1): 25–44.

Bates, M. 1923. A Study of the Müller-Lyer Illusion with Special Reference to Paradoxical Movement and the Effect of Attitude. *American Journal of Psychology* 34: 46–72.

Block, N. 2010. Attention and Mental Paint. *Philosophical Issues* 20: 23–63.

[30] For a critique in similar spirit concerning colour illusions see Kalderon 2011.

[31] I am very grateful to audiences at Sheffield and KCL for their questions and discussion, and especially to my commentator on the latter occasion, Jørgen Dyrstad, and to Nick Shea, Matt Parrott and David Owens for extended discussion. Thanks also to Boyd Millar for feedback and especially to Craig French and Chris Hill for extremely helpful written comments. Thanks, finally, to Hanna Pickard for her advice and encouragement.

Brewer, B. 2004. Realism and the Nature of Perceptual Experience. *Philosophical Issues* 14(1): 61–77.

Brewer, B. 2008. How to Account for Illusion. In A. Haddock and F. Macpherson (eds) *Disjunctivism: Perception, Action, Knowledge*. Oxford: Oxford University Press, pp. 168–180.

Brewer, B. 2011. *Perception and Its Objects*. Oxford: Oxford University Press.

Brewer, B. 2013. Attention and Direct Realism. *Analytic Philosophy* 54(4): 421–435.

Brown, L. B. and Houssiadas, L. 1965. The Perception of Illusions as a Constancy Phenomenon. *British Journal of Psychology* 56(2/3): 135–141.

Burge, T. 2005. Disjunctivism and Perceptual Psychology. *Philosophical Topics* 33(1): 1–78.

Byrne, A. 2009. Experience and Content. *Philosophical Quarterly* 59(236): 429–451.

Campbell, J. 2002. *Reference and Consciousness*. Oxford: Oxford University Press.

Campbell, J. 2009. Consciousness and Reference. In B. McLaughlin, A. Beckermann and S. Walter (eds) *Oxford Handbook of Philosophy of Mind*. Oxford: Oxford University Press.

Campbell, J. 2010. Demonstrative Reference, the Relational View of Experience and the Proximality Principle. In R. Jeshion (ed.) *New Essays on Singular Thought.* Oxford: Oxford University Press.

Craig, E. 1976. Sensory Experience and the Foundations of Knowledge. *Synthese* 33(1): 1–24.

Day, R. H. 1962. The Effects of Repeated Trials and Prolonged Fixation on Error in the Müller-Lyer Figure. *Psychological Monographs* 76(14): 1–19.

Day, R. H. 1965. Inappropriate Constancy Explanation of Spatial Distortions. *Nature* 207: 891–893.

Day, R. H. 1989. Natural and Artificial Cues, Perceptual Compromise and the Basis of Veridical and Illusory Perception. In D. Vickers and P. L. Smith (eds) *Human Information Processing: Measures, Mechanisms and Models*. California: North Holland, pp. 107–129.

Delboeuf, J. 1892. Sur une nouvelle illusion d'optique. *Bulletin de l'Académie Royale de Belgique* 24(3d serie): 545–558.

DeLucia, P. R. and Hochberg J. 1991. Geometrical Illusions in Solid Objects Under Ordinary Viewing Conditions. *Perception & Psychophysics* 50(6): 547–54.

Eaglen, J. and Kirkwood, B. 1970. The Effect of Instructions on Judgment of the Müller-Lyer Illusion with Normal and Haptically Mediated Visual Inspection. *Perception & Psychophysics* 8(1): 35–36.

Erlebacher, A. and Sekuler, R. 1969. Explanation of the Müller-Lyer Illusion: Confusion Theory Examined. *Journal of Experimental Psychology*. 80(3): 462–467.

Fish, W. 2009. *Perception, Hallucination, and Illusion*. Oxford: Oxford University Press.

Foster, J. 2000. *The Nature of Perception*. Oxford: Oxford University Press.

French, C. Forthcoming. Idiosyncratic Perception. *The Philosophical Quarterly* doi:10.1093/pq/pqv068

Gardner, R. W. and Long, R. I. 1961. Selective Attention and the Müller-Lyer Illusion. *Psychological Record* 11: 317–320.

Genone, J. 2014. Appearance and Illusion. *Mind* 123(490): 339–376.

Green, D. M. and Swets, J. A. 1966. *Signal detection theory and psychophysics*. New York: Wiley.

Green, F. A. 1972. *The Relationship between Visual Illusions and Cues to Distance*. Durham theses, Durham University. Available at Durham E-Theses Online: http://etheses.dur.ac.uk/9323/

Gregory, R. L. 1963. Distortion of Visual Space as Inappropriate Constancy Scaling. *Nature* 199: 678–680.

Gregory, R. L. 1964. Illusory Perception as a Constancy Phenomenon. *Nature* 204: 302–3.

Grove, P. M., Ashton, J., Kawachi, Y. and Sakurai, K. 2012. Auditory Transients Do Not Affect Visual Sensitivity in Discriminating Between Objective Streaming and Bouncing Events. *Journal of Vision* 12(8): 1–11.

Holding, D. H. 1970. A Line Illusion with Irrelevant Depth Cues. *American Journal of Psychology* 83: 280–282.

Howe, C. Q. and Purves, P. 2005. The Müller-Lyer Illusion Explained by the Statistics of Image-Source Relationships. *Proceedings of the National Academy of Sciences of the United States of America* 102: 1234–1239.

Humphrey, N. K. and Morgan, M. J. 1965. Constancy and the Geometric Illusions. *Nature* 206: 744–45.

Hyman, J. 2006. *The Objective Eye: colour, form and reality in the theory of art*. The University of Chicago.

Jansari, A., Miller, S., Pearce, L., Cobb, S., Sagiv, N., Williams, A., Tree, J., and Hanley R. 2015. The Man Who Mistook His Neuropsychologist for a Popstar: When Configural Processing Fails in Acquired Prosopagnosia. *Frontiers in Human Neuroscience* 9(390): 1–16.

Kalderon, M. E. 2011. Color Illusion. *Noûs* 45(4): 751–775.

Macmillan, N. A. and Creelman, D. 2005. *Detection theory: A user's guide* (2nd ed.). Mahwah, NJ: Erlbaum.

Martin, M. G. F. 2002. The Transparency of Experience. *Mind and Language* 4(4): 376–425.

Martin, M. G. F. 2004. The Limits of Self-Awareness. *Philosophical Studies* 120: 37–89.

Martin, M. G. F. 2006. On Being Alienated. In T. Gendler and J. Hawthorne (eds) *Perceptual Experience* Oxford: Oxford University Press, pp. 354–410.

Martin, M. G. F. 2009 [1997]. The Reality of Appearances. Reprinted in H. Logue and A. Byrne (eds) *Disjunctivism: Contemporary Readings*. MIT Press, pp. 91–115.

Martin, M. G. F. 2010. What's in a Look? In B. Nanay (ed.) *Perceiving the World*. Oxford: Oxford University Press, pp. 160–225.

McLaughlin, B. P. 2010. The Representational Vs. The Relational View of Visual Experience. *Royal Institute of Philosophy Supplement* 85(67): 239–262.

Millar, B. 2015. Naïve Realism and Illusion. *Ergo* 2: 607–625.

Moore, G. E. 1953. Sense-data. In his *Some Main Problems of Philosophy*. London: George Allen & Unwin, pp. 28–40.

Morgan, M. J. and Glennerster, A. 1991. Efficiency of Locating Centres of Dot-Clusters by Human Observers. *Vision Research* 31: 2075–2083.

Morgan, M. J., Dillenburger, B., Raphael, S. and Solomon, J. A. 2012. Observers can Voluntarily Shift Their Psychometric Functions without Losing Sensitivity. *Attention, Perception, & Psychophysics* 74: 185–193.

Morgan, M. J., Hole, G. J. and Glennerster, A. 1990. Biases and Sensitivities in Geometric Illusions. *Vision Research* 30: 1793–1810.

Mukerji, N. 1957. Is "illusion" illusion? *Journal of General Psychology* 57: 209–212.

Müller-Lyer, F. C. 1889. Optische Urteilstauschungen. *Dubois-Reymond's Archive für Anatomie und Physiologie, Physiologische Abteilung* 2: 261 – 267.

Mundy, M. E. 2014. Testing day: The effects of Processing Bias Induced by Navon Stimuli on the Strength of the Müller-Lyer Illusion. *Advanced in Cognitive Psychology* 10(1): 9–14.

Navon, D. 1977. Forest Before the Trees: The precedence of Global Features in Visual Perception. *Cognitive Psychology* 9: 353–383.

Ninio, J. 2014. Geometrical illusions are not always where you think they are: a review of some classical and less classical illusions, and ways to describe them. *Frontiers in Human Neuroscience* 8(856): 1–21.

Noë, A. 2004. *Action in Perception*. Cambridge, MA: MIT Press.

Palmer, S. E. 1999. *Vision Science.* Cambridge, MA: MIT Press.

Peterson, W. W., Birdsall, T. G. and Fox, W. C. 1954. The Theory of Signal Detectability. *Transactions of the IRE Professional Group on Information Theory* PGIT-4: 171–212.

Phillips, I. B. 2005. *Experience and Intentional Content*. BPhil Thesis. Oxford University.

Pressey, A. and Martin, N. S. 1990. The Effects of Varying Fins in Müller-Lyer and Holding Illusions. *Psychological Research* 52(1): 46–53

Pressey, A. W. 1967. A theory of the Mueller-Lyer illusion. *Perceptual and Motor Skills* 25: 569–572.

Pressey, A. W. 1971. An Extension of Assimilation Theory to Illusions of Size, Area, and Direction. *Perception & Psychophysics* 9: 172–176.

Pressey, A. W. 2013. Gregory's Theory of Visual Illusion: A case study of wilful blindness. *Pressey's Pieces*, www.awpressey.com

Price, H. H. 1932. *Perception.* London: Methuen & Co.

Prinzmetal, W. Shimamura, A. P. and Mikolinski, M. 2001. The Ponzo Illusion and the Perception of Orientation. *Perception & Psychophysics* 63(1): 99–114.

Rock, I. 1984. *Perception*. New York: Scientific American Library.

Russell, B. 1912. *The Problems of Philosophy*. Oxford: Oxford University Press.

Sadza, K. J. and de Weert, C. M. M. 1984. Influence of color and luminance on the Muller-Lyer illusion. Perception &Psychophysics 35(3): 214–220.

Sanford, E. C. 1897. *Experimental Psychology*. London: D. C. Heath & Co.

Searle, S. 2015. *Seeing Things as They Are: A Theory of Perception*. Oxford University Press.

Sekuler, R. and Erlebacher, A. 1971. The Two Illusions of Müller-Lyer: Confusion Theory Reexamined. *American Journal of Psychology* 84(4): 477–86.

Sekuler, R., Sekuler, A. B. and Lau, R. 1997. Sound Alters Visual Motion Perception. *Nature* 385(6614): 308.

Siegel, S. 2011. *The Contents of Visual Experience*. New York: Oxford University Press.

Smith, A. D. 2010. Disjunctivism and Illusion. *Philosophy and Phenomenological Research* 80(2): 384–410.

Swets, J. A., Tanner, W. P., Jr. and Birdsall, T. G. 1961. Decision Processes in Perception. *Psychological Review* 68: 301–340.

Tausch, R. 1954. Optische Taüschungen als artifizielle Effekte der Gestaltungs-prozesse von Grössen und Formenkonstanz in der natürlichen Raumwahrnehmung. *Psychologische Forschung* 24: 299–348.

Thiéry, A. 1896. Uber Geometrisch-Optische Taüschungen. *Philosophische Studien* 12: 67–126.

Travis, C. 2004. The Silence of the Senses. *Mind* 113(449): 57–94.

Tye, M. 2014. What is the Content of a Hallucinatory Experience? In B. Brogaard (ed.) *Does Perception have Content?* Oxford: Oxford University Press, pp. 291–308.

Waite, H. and Massaro, D. W. 1970. A Test of Gregory's Constancy-Scaling Explanation of the Müller-Lyer illusion. *Nature* 227: 733–734.

Williamson, T. forthcoming. Model-Building in Philosophy. In R. Blackford and D. Broderick (eds) *Philosophy's Future: The Problem of Philosophical Progress*. Oxford: Wiley.

Witt, J. K., Taylor, J. E. T., Sugovic, M. and Wixted, J. T. 2015. Signal Detection Measures Cannot Distinguish Perceptual Biases From Response Biases. *Perception* 44: 289–300.

Woloszyn, M. R. 2010. Contrasting Three Popular Explanations for the Müller-Lyer Illusion. *Current Research in Psychology* 1(2): 102–107.

Woodworth, R. S. 1938. *Experimental psychology*. New York: Holt.